

Published Date	16 June 2025	Portal	Analytics India Magazine
Edition	Online	Category	A
Headline	L&T Finance Hopes Project Cyclops Would Be 90% Agentic One Day		

**L&T Finance Hopes Project Cyclops Would Be 90% Agentic One Day**

<https://analyticsindiamag.com/ai-features/lt-finance-hopes-project-cyclops-would-be-90-agentic-one-day/>

In the world of finance, the use of AI is tricky owing to privacy issues, the dreaded risk of hallucinations and the guardrails. Given that modern systems are increasingly getting foolproof, there is a huge opportunity for changing the industry with agents—all with a little fine-tuning.

Debarag Banerjee, chief AI and data officer at L&T Finance, spoke with AIM about how the firm is steadily moving away from rules-based automation towards a future led by agents.

The most important and real transformation happening under the hood is with Project Cyclops, L&T Finance’s proprietary AI stack for real-time, high-accuracy credit decisioning.

“We launched this with our two-wheeler portfolio last year. It now handles 100% of those loans,” Banerjee said. “This year, we’ve extended it to our tractor business and are preparing to roll it out for small business loans.”

Project Cyclops pulls together various “trust signals”—from customer profiles to repayment behaviour—into an ensemble model that can instantly separate delinquent-risk borrowers from credit-worthy ones. “You upload your information and, just like that, you get a decision,” Banerjee said.

**How to Build This in India?**

In terms of data residency and compliance, L&T Finance is already future-proofing things. “Even for closed-source LLMs, we insist on endpoints hosted in India to prepare for laws like the Digital Personal Data Protection Act (DPDPA),” Banerjee added.

Open-source models naturally provide better control. “Since we host and manage the stack ourselves, the data is more secure. We ensure contractually that our EII data isn’t used for model retraining.” Much of the data used to fine-tune open-source LLMs is proprietary. Even when it’s not proprietary, the formulation under which the team trains them makes the contextual usage proprietary.

L&T Finance has also explored Indian LLMs built for Indic languages. “They’re a good start. But the number of parameters still matters.”

Meanwhile, Project Cyclops, the firm’s proprietary ML stack, continues to scale. It combines models across various trust signals — from customer data to repayment behaviour — and re-ensembles them to deliver real-time credit decisions.

The company deliberately took a multi-LLM route from day one. The goal is flexibility, not being locked into any single provider or model.

“Instead of being tied to LLMs from any one company, our stack can call any model our developer thinks is right for that task,” said Banerjee. This includes Google’s Gemini (multiple versions), OpenAI models through Azure, and several open-source LLMs hosted on GPU-as-a-service platforms.

They’ve also tested Meta’s Llama family (3.1, 3.2) and successfully fine-tuned them for performance comparable to larger models like Gemini, but with lower inference costs.

“In one of our other applications, we found medium-sized fine-tuned Llamas performing nearly as well as some of the premium models,” he noted. “We’re agnostic to geography or company, as long as data privacy is maintained and we retain full control.”

### **Tackling Bias and Ethics in AI Lending**

When asked about ethical concerns in using AI for credit decisions, especially cases where background visuals or personal environment might influence model behaviour, Banerjee stressed two things: statistical validation and consent.

“Any trust signal we use has to hold up against statistically significant past data. If it’s frivolous, it gets discarded,” he said. “We are also very careful about consent. All data usage is fully transparent to the customer.”

He acknowledged the risk of adversarial behaviour, like customers gaming the system with artificial backgrounds. “But these kinds of patterns are caught through quality checks and operational safeguards,” Banerjee noted. “It’s a team effort — credit, risk, field ops, everyone must align for the system to work at scale.”

Traditional software relied on endless rules. With agent frameworks powered by LLMs and task-specific tools, Banerjee said that the effort to build systems has reduced drastically. “You can create something with a minimalist approach, deploy it, and let it improve over time. These are not just tools — they’re self-improving systems,” Banerjee said.

The future of agentic AI is both exciting and inevitable for Banerjee. “Agentic AI seems to have finally caught that right gap,” he observed. “We can already see it proving its mettle—not only in decision-making but in emulating human functions.”

He envisions agents becoming self-improving and minimalist in design. “Instead of writing rule after rule, you get to a working solution quickly, test it, improve it, and even reinforce it.”

Though we’re not fully there yet, the progress is palpable. “I’ve seen demos where agents can do EDA, generate train-test sets, build models, and write documentation—all with very little human intervention,” he said. “That future, where something like Project Cyclops is 90% agent-driven, would be wonderful.”

### **Regulation, Black Swans & the Human Touch**

Will agents eventually monitor finance and trading platforms on their own? “There are still regulatory needs—maker, checker, monitor—which will stay,” he said. “And while you can create agents for predictable failures, black swans are by nature unpredictable.”

He also believes AI will create jobs. “One area where AI may generate jobs is in humans playing both white hat and black hat—looking for ways AI can fail or be misused and figuring out how to recover.” For L&T Finance, agentic AI is not just about tech, it’s about solving for India’s underserved.

“I was there when India got connected. Then came digital payments. Now, the next big inflection is digital access to credit for the bottom and middle of the pyramid,” he said. “India has the opportunity to leapfrog old credit systems because its consumers are digitally connected.”



In the world of finance, the use of AI is tricky owing to privacy issues, the dreaded risk of hallucinations and the guardrails. Given that modern systems are increasingly getting foolproof, there is a huge opportunity for changing the industry with agents—all with a little fine-tuning.

[Debaraj Banerjee](#), chief AI and data officer at L&T Finance, spoke with **AIM** about how the firm is steadily moving away from rules-based automation towards a future led by agents.

The most important and real transformation happening under the hood is with Project Cyclops, L&T Finance's proprietary AI stack for real-time, high-accuracy credit decisioning.

"We launched this with our two-wheeler portfolio last year. It now handles 100% of those loans," Banerjee said. "This year, we've extended it to our tractor business and are preparing to roll it out for small business loans."

[Project Cyclops](#) pulls together various "trust signals"—from customer profiles to repayment behaviour—into an ensemble model that can instantly separate delinquent-risk borrowers from credit-worthy ones. "You upload your information and, just like that, you get a decision," Banerjee said.

## How to Build This in India?

In terms of data residency and compliance, L&T Finance is already future-proofing things. "Even for closed-source LLMs, we insist on endpoints hosted in India to prepare for laws like the Digital Personal Data Protection Act (DPDPA)," Banerjee added.

Open-source models naturally provide better control. "Since we host and manage the stack ourselves, the data is more secure. We ensure contractually that our EII data isn't used for model retraining." Much of the data used to fine-tune open-source LLMs is proprietary. Even when it's not proprietary, the formulation under which the team trains them makes the contextual usage proprietary.

L&T Finance has also explored Indian LLMs built for Indic languages. "They're a good start. But the number of parameters still matters."

Meanwhile, Project Cyclops, the firm's proprietary ML stack, continues to scale. It combines models across various trust signals — from customer data to repayment behaviour — and re-ensembles them to deliver real-time credit decisions.

The company deliberately took a multi-LLM route from day one. The goal is flexibility, not being locked into any single provider or model.

"Instead of being tied to LLMs from any one company, our stack can call any model our developer thinks is right for that task," said Banerjee. This includes Google's Gemini (multiple versions), OpenAI models through Azure, and several open-source LLMs hosted on GPU-as-a-service platforms.

They've also tested Meta's [Llama](#) family (3.1, 3.2) and successfully fine-tuned them for performance comparable to larger models like Gemini, but with lower inference costs.

"In one of our other applications, we found medium-sized fine-tuned Llamas performing nearly as well as some of the premium models," he noted. "We're agnostic to geography or company, as long as data privacy is maintained and we retain full control."

## Tackling Bias and Ethics in AI Lending

When asked about ethical concerns in using AI for credit decisions, especially cases where background visuals or personal environment might influence model behaviour, Banerjee stressed two things: statistical validation and consent.

"Any trust signal we use has to hold up against statistically significant past data. If it's frivolous, it gets discarded," he said. "We are also very careful about consent. All data usage is fully transparent to the customer."

He acknowledged the risk of adversarial behaviour, like customers gaming the system with artificial backgrounds. "But these kinds of patterns are caught through quality checks and operational safeguards," Banerjee noted. "It's a team effort — credit, risk, field ops, everyone must align for the system to work at scale."

Traditional software relied on endless rules. With agent frameworks powered by LLMs and task-specific tools, Banerjee said that the effort to build systems has reduced drastically. "You can create something with a minimalist approach, deploy it, and let it improve over time. These are not just tools — they're self-improving systems," Banerjee said.

The future of agentic AI is both exciting and inevitable for Banerjee. "Agentic AI seems to have finally caught that right gap," he observed. "We can already see it proving its mettle—not only in decision-making but in emulating human functions."

He envisions agents becoming self-improving and minimalist in design. "Instead of writing rule after rule, you get to a working solution quickly, test it, improve it, and even reinforce it."

Though we're not fully there yet, the progress is palpable. "I've seen demos where agents can do EDA, generate train-test sets, build models, and write documentation—all with very little human intervention," he said. "That future, where something like Project Cyclops is 90% agent-driven, would be wonderful."

### **Regulation, Black Swans & the Human Touch**

Will agents eventually monitor finance and trading platforms on their own? "There are still regulatory needs—maker, checker, monitor—which will stay," he said. "And while you can create agents for predictable failures, black swans are by nature unpredictable."

He also believes AI will create jobs. "One area where AI may generate jobs is in humans playing both white hat and black hat—looking for ways AI can fail or be misused and figuring out how to recover." For L&T Finance, agentic AI is not just about tech, it's about solving for India's underserved.

"I was there when India got connected. Then came digital payments. Now, the next big inflection is digital access to credit for the bottom and middle of the pyramid," he said. "India has the opportunity to leapfrog old credit systems because its consumers are digitally connected."